

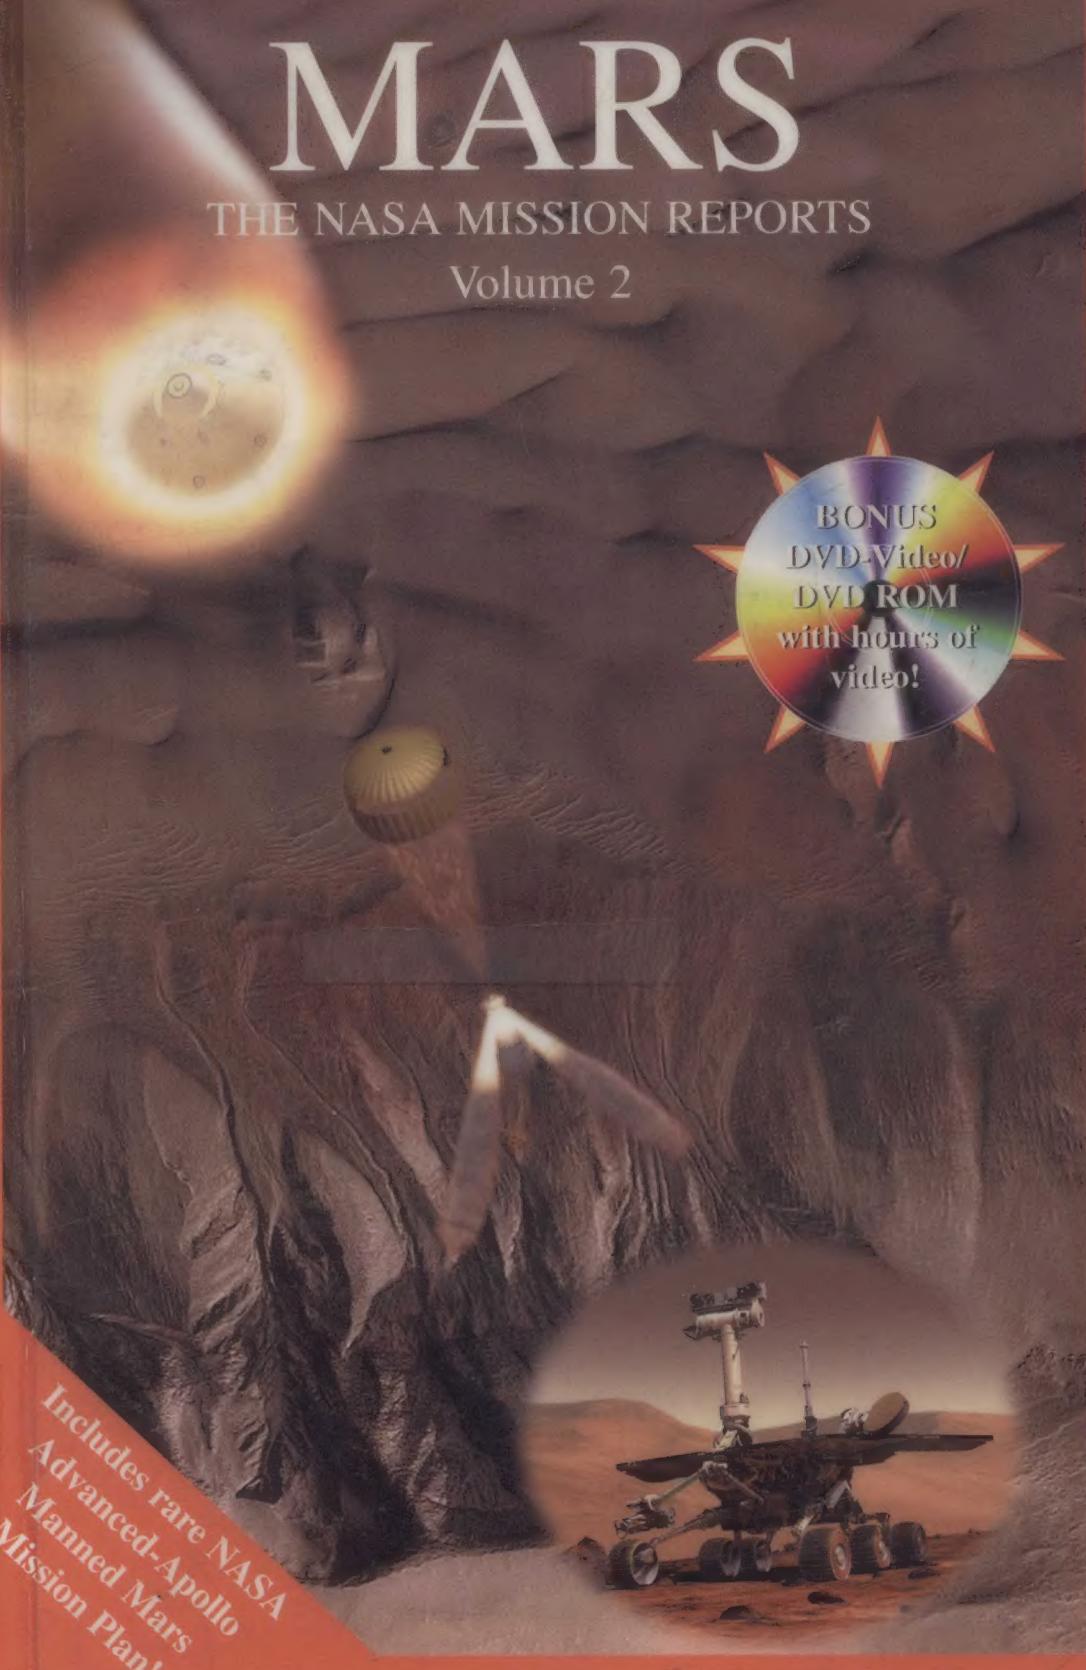
MARS

THE NASA MISSION REPORTS

Volume 2



BONUS
DVD-Video/
DVD ROM
with hours of
video!



Includes rare NASA
Advanced-Apollo
Manned Mars
Mission Plan!

3 7244 1624 8176 7

In 1877 the famed Italian astronomer Giovanni brand-new 8.6 inch telescope to study the planets. To his great surprise he suspected that he saw symmetry on Mars. In the years that followed one astronomer after another looked at the red planet and gradually a mythology was formed—a mythology of alien intellect.

By the 1890's the martial influence had spilled over into all walks of life and sparked philosophical debates and wonderous fictions. Scientists, fantasists and people of all creeds looked up and wondered—is there life out there?

Now, more than a century later, nations around the world are bombarding Mars with an unprecedented fleet of exploratory vehicles. Their journey taking less time than it took Amundsen and Shackleton to reach the poles of Earth, these small but hardy robotic emissaries are thrusting their way through the depths of interplanetary space to take up residence in the barren Martian deserts. Their goal is to answer one of the oldest questions in mankind's history. Is there life out there?

In this sequel to the best-selling first volume, the reader is brought up to date with the most recent results from our nearest neighbour. Filled with a wealth of facts about the latest fleet of Martian explorers as well as a look at what may be coming next in mankind's most ambitious quest for knowledge.

Reviews of Mars - The NASA Mission Reports.

"Godwin has assembled an authoritative, blow-by-blow resource for serious space buffs, and it's this breadth that makes the book a standout even in this excellent series." **Amazon.com editorial**

"A well-conceived book which those interested in planetary exploration cannot afford to miss." **Spaceflight - British Interplanetary Society**

"Mars buffs may find the CD-ROM alone to be worth the cover price." **Astronomy**

"This is an outstanding reference at a reasonable price." **SkyNews**

"The book and CD-ROM have a lot to recommend them, particularly given the low price, and would form a valuable addition to the bookshelf of anybody interested in Mars exploration."

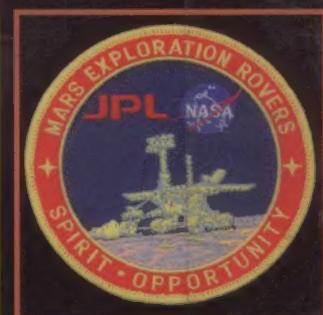
The Observatory

DVD-Video* includes: The Landing of Opportunity, JPL Opportunity Press Conference, MER Animation**, M2K4 Animation, MER Launch & Preflight
Exclusive Interviews with: MER Athena Science Payload Principal Investigator Steve Squyres, MER Entry Descent & Landing Manager Rob Manning and Opportunity Mission Manager Jim Erickson.

Also includes:

A rare 1975 lecture by Dr Wernher von Braun discussing missions to Mars.
DVD-ROM includes NASA publications about Mars: The Book of Mars, Humans to Mars, The Martian Landscape, Study of Life Support Systems for Space Missions Exceeding 1 year in Duration, Report of 90 Day Study for the Moon & Mars
PLUS! Mars Exploration Rover Imagery up to Spirit Sol 47 & Opportunity Sol 29, Odyssey THEMIS Imagery, JPL Opportunity Meridiani Planum Water Press Conference March 2nd 2004, Canadian Space Agency Director Marc Garneau speech on Canadian Mars plans.

* NTSC Region 0, ** Courtesy MAAS Digital/Cornell



ISBN 1-894959-05-1



9 781894 959056



Box 62034, Burlington
Ontario, L7R 4K2, Canada
<http://www.apogeebooks.com>

US \$28.95
Can \$38.95
UK £20.95

"We see the twin rovers as stepping stones for the rest of the decade and to a future decade of Mars exploration that will ultimately provide the knowledge necessary for human exploration," said Orlando Figueroa, director of the Mars Exploration Program at NASA Headquarters.

JPL, a division of the California Institute of Technology in Pasadena, manages the Mars Exploration Rover project for NASA's Office of Space Science, Washington.

For information about the Mars Exploration Rover project on the Internet, visit:
<http://mars.jpl.nasa.gov/mer>

NASA will feature live webcasts of the launches on the Internet at:
<http://www.jpl.nasa.gov/webcast/mer>

Cornell University's web site on the science payload is at:
<http://athena.cornell.edu>

Media Services Information

NASA Television Transmission

NASA Television is broadcast on the satellite AMC-2, transponder 9C, C band, 85 degrees west longitude, frequency 3880.0 MHz, vertical polarization, audio monaural at 6.8 MHz. The schedule for Mars arrival television transmissions will be available from the Jet Propulsion Laboratory, Pasadena, California; and NASA Headquarters, Washington.

Launch Media Credentialing

News media representatives who would like to cover the launch in person must be accredited through the NASA Kennedy Space Center newsroom. Journalists may contact the newsroom at 321/867-2468 for more information.

Briefings

An extensive schedule of news and background briefings will be held at JPL during the landing period, with later briefings originating jointly from JPL and NASA Headquarters. A schedule of briefings is available on the Internet at JPL's Mars News site (below).

Internet Information

Extensive information on the Mars Exploration Rover project including an electronic copy of this press kit, press releases, fact sheets, status reports, briefing schedule and images, is available from the Jet Propulsion Laboratory's Mars Exploration Rover newsroom website: <http://www.jpl.nasa.gov/mer>. The Mars Exploration Rover project also maintains a web site at: <http://mars.jpl.nasa.gov/mer>. Cornell University's web site on the science payload is at: <http://athena.cornell.edu>.

Quick Facts

Spacecraft

Cruise vehicle dimensions: 2.65 meters (8.7 feet) diameter, 1.6 meters (5.2 feet) tall
Rover dimensions: 1.5 meter (4.9 feet) high by 2.3 meters (7.5 feet) wide by 1.6 meter (5.2 feet) long
Weight: 1,062 kilograms (2,341 pounds) total at launch, consisting of 174-kilogram (384-pound) rover, 365-kilogram (805-pound) lander, 198-kilogram (436-pound) backshell and parachute, 90-kilogram (198-pound) heat shield and 183-kilogram (403-pound) cruise stage, plus 52 kilograms (115 pounds) of propellant
Power: Solar panel and lithium-ion battery system providing 140 watts on Mars surface
Science instruments: Panoramic cameras, miniature thermal emission spectrometer, Mössbauer spectrometer, alpha particle X-ray spectrometer, microscopic imager, rock abrasion tool, magnet arrays

Rover A Mission

Launch vehicle: Delta II 7925
Launch period: June 8-24, 2003
Earth-Mars distance at launch: 105 million kilometers (65 million miles)
Mars landing: January 4, 2004, at about 2 p.m. local Mars time (8:11 p.m. January 3 PST)
Landing site: Gusev Crater, possible former lake in giant impact crater
Earth-Mars distance on landing day: 170.2 million kilometers (105.7 million miles)
One-way speed-of-light time Mars-to-Earth on landing day: 9.46 minutes
Total distance traveled Earth to Mars (approximate): 500 million kilometers (311 million miles)
Near-surface atmospheric temperature at landing site: -100 C (-148 F) to 0 C (32 F)
Primary mission: 90 Mars days, or "sols" (equivalent to 92 Earth days)

NASA will feature live webcasts of the launches on the Internet at:

<http://www.jpl.nasa.gov/webcast/mer>

Cornell University's web site on the science payload is at:

<http://athena.cornell.edu>

Media Services Information

NASA Television Transmission

NASA Television is broadcast on the satellite AMC-2, transponder 9C, C band, 85 degrees west longitude, frequency 3880.0 MHz, vertical polarization, audio monaural at 6.8 MHz. The schedule for Mars arrival television transmissions will be available from the Jet Propulsion Laboratory, Pasadena, California; and NASA Headquarters, Washington.

Launch Media Credentialing

News media representatives who would like to cover the launch in person must be accredited through the NASA Kennedy Space Center newsroom. Journalists may contact the newsroom at 321/867-2468 for more information.

Briefings

An extensive schedule of news and background briefings will be held at JPL during the landing period, with later briefings originating jointly from JPL and NASA Headquarters. A schedule of briefings is available on the Internet at JPL's Mars News site (below).

Internet Information

Extensive information on the Mars Exploration Rover project including an electronic copy of this press kit, press releases, fact sheets, status reports, briefing schedule and images, is available from the Jet Propulsion Laboratory's Mars Exploration Rover newsroom website: <http://www.jpl.nasa.gov/mer>. The Mars Exploration Rover project also maintains a web site at: <http://mars.jpl.nasa.gov/mer>. Cornell University's web site on the science payload is at: <http://athena.cornell.edu>.

Quick Facts

Spacecraft

Cruise vehicle dimensions: 2.65 meters (8.7 feet) diameter, 1.6 meters (5.2 feet) tall

Rover dimensions: 1.5 meter (4.9 feet) high by 2.3 meters (7.5 feet) wide by 1.6 meter (5.2 feet) long

Weight: 1,062 kilograms (2,341 pounds) total at launch, consisting of 174-kilogram (384-pound) rover, 365-kilogram (805-pound) lander, 198-kilogram (436-pound) backshell and parachute, 90-kilogram (198-pound) heat shield and 183-kilogram (403-pound) cruise stage, plus 52 kilograms (115 pounds) of propellant

Power: Solar panel and lithium-ion battery system providing 140 watts on Mars surface

Science instruments: Panoramic cameras, miniature thermal emission spectrometer, Mössbauer spectrometer, alpha-particle X-ray spectrometer, microscopic imager, rock abrasion tool

Media Contacts

Donald Savage
Headquarters Washington, D.C.

Policy / Program Management

202/358-1547
donald.savage@hq.nasa.gov

Guy Webster
Jet Propulsion Laboratory, Pasadena, California

Mars Exploration Rover Mission

818/354-5011
guy.webster@jpl.nasa.gov

David Brand
Cornell University, Ithaca, New York

Science Payload

607/255-3651
deb27@cornell.edu

George Diller
Kennedy Space Center, Florida

Launch

321/867-2468
george.diller-1@ksc.nasa.gov

GENERAL RELEASE: NASA PREPARES TWO ROBOT ROVERS FOR MARS EXPLORATION

NASA's Mars Exploration Rover project kicks off by launching the first of two unique robotic geologists, as early as June 8. The identical rolling rovers see sharper images, can explore farther and examine rocks better than anything that's ever landed on Mars. The second rover mission, bound for a different site on Mars, will launch as soon as June 25.

"The instrumentation onboard these rovers, combined with their great mobility, will offer a totally new view of Mars, including a microscopic view inside rocks for the first time," said Dr. Ed Weiler, associate administrator for space science, NASA Headquarters, Washington.

"However, missions to Mars have proven to be far more hazardous than missions to other planets. Historically, two out of three missions, from all countries who have tried to land on Mars, ended in failure. We have done everything we can to ensure our rovers have the best chance of success, and today I gave the order to proceed to launch," Weiler said.

The first rover will arrive at Mars on January 4, 2004, the second on January 25. Plans call for each to operate for at least three months. These missions continue NASA's quest to understand the role of water on Mars. "We will be using the rovers to find rocks and soils that could hold clues about wet environments of Mars' past," said Dr. Cathy Weitz, Mars Exploration Rover program scientist at NASA Headquarters. "We'll analyze the clues to assess whether those environments may have been conducive to life."

First, the rovers have to safely reach Mars. "The rovers will use innovations to aid in safe landings, but risks remain," said Peter Theisinger, Mars Exploration rover project manager at NASA's Jet Propulsion Laboratory, Pasadena, California.

The rovers will bounce to airbag-cushioned landings at sites offering a balance of favorable conditions for safe landings and interesting science. The designated site for the first mission is Gusev Crater. The second rover will go to a site called Meridiani Planum. "Gusev and Meridiani give us two different types of evidence about liquid water in Mars' history," said Dr. Joy Crisp, Mars Exploration Rover project scientist at JPL. "Gusev appears to have been a crater lake. The channel of an ancient riverbed indicates water flowed right into it. Meridiani has a large deposit of gray hematite, a mineral that usually forms in a wet environment," Crisp said.

The rovers, working as robotic field geologists, will examine the sites for clues about what happened there. "The clues are in the rocks, but you can't go to every rock, so you split the job into two pieces," said Dr. Steve Squyres of Cornell University, Ithaca, N.Y., principal investigator for the package of science instruments on the rovers.

First, a panoramic camera at human-eye height, and a miniature thermal emission spectrometer, with infrared vision, help scientists identify the most interesting rocks. The rovers can watch for hazards in their way and maneuver around them. Each six-wheeled robot has a deck of solar panels, about the size of a kitchen table, for power. The rover drives to the selected rock and extends an arm with tools on the end. Then, a microscopic imager, like a geologist's hand lens, gives a close-up view of the rock's texture. Two spectrometers identify the composition of the rock. The fourth tool substitutes for a geologist's hammer. It exposes the fresh interior of a rock by scraping it away the weathered surface layer.

Both rover missions will lift off from Cape Canaveral Air Force Station, Florida, on Delta II launch vehicles. Launch opportunities begin for the first mission at 2:06 p.m. EDT June 8 and for the second mission at 12:38 a.m. EDT June 25, and repeat twice daily for up to 21 days for each mission.

"We see the twin rovers as stepping stones for the rest of the decade and to a future decade of Mars exploration that will ultimately provide the knowledge necessary for human exploration," said Orlando Figueroa, director of the Mars Exploration Program at NASA Headquarters.

JPL, a division of the California Institute of Technology in Pasadena, manages the Mars Exploration Rover project for NASA's Office of Space Science, Washington.

For information about the Mars Exploration Rover project on the Internet, visit:
<http://mars.jpl.nasa.gov/mer>

NASA will feature live webcasts of the launches on the Internet at:
<http://www.jpl.nasa.gov/webcast/mer>

Cornell University's web site on the science payload is at:
<http://athena.cornell.edu>

Media Services Information

NASA Television Transmission

NASA Television is broadcast on the satellite AMC-2, transponder 9C, C band, 85 degrees west longitude, frequency 3880.0 MHz, vertical polarization, audio monaural at 6.8 MHz. The schedule for Mars arrival television transmissions will be available from the Jet Propulsion Laboratory, Pasadena, California; and NASA Headquarters, Washington.

Launch Media Credentialing

News media representatives who would like to cover the launch in person must be accredited through the NASA Kennedy Space Center newsroom. Journalists may contact the newsroom at 321/867-2468 for more information.

Briefings

An extensive schedule of news and background briefings will be held at JPL during the landing period, with later briefings originating jointly from JPL and NASA Headquarters. A schedule of briefings is available on the Internet at JPL's Mars News site (below).

Internet Information

Extensive information on the Mars Exploration Rover project including an electronic copy of this press kit, press releases, fact sheets, status reports, briefing schedule and images, is available from the Jet Propulsion Laboratory's Mars Exploration Rover newsroom website: <http://www.jpl.nasa.gov/mer>. The Mars Exploration Rover project also maintains a web site at: <http://mars.jpl.nasa.gov/mer>. Cornell University's web site on the science payload is at: <http://athena.cornell.edu>.

Spacecraft

Cruise vehicle dimensions: 2.65 meters (8.7 feet) diameter, 1.6 meters (5.2 feet) tall
Rover dimensions: 1.5 meter (4.9 feet) high by 2.3 meters (7.5 feet) wide by 1.6 meter (5.2 feet) long
Weight: 1,062 kilograms (2,341 pounds) total at launch, consisting of 174-kilogram (384-pound) rover, 365-kilogram (805-pound) lander, 198-kilogram (436-pound) backshell and parachute, 90-kilogram (198-pound) heat shield and 183-kilogram (403-pound) cruise stage, plus 52 kilograms (115 pounds) of propellant

Power: Solar panel and lithium-ion battery system providing 140 watts on Mars surface
Science instruments: Panoramic cameras, miniature thermal emission spectrometer, Mössbauer spectrometer, alpha particle X-ray spectrometer, microscopic imager, rock abrasion tool, magnet arrays

Rover A Mission

Launch vehicle: Delta II 7925
Launch period: June 8-24, 2003
Launch distance at launch: 105 million kilometers (65 million miles)
Earth-Mars distance at launch: about 2 p.m. local Mars time (8:11 p.m. January 3 PST)

Mars landing: January 4, 2004, at about 2 p.m. local Mars time (8:11 p.m. January 3 PST)
Landing site: Gusev Crater, possible former lake in giant impact crater
Landing distance on landing day: 170.2 million kilometers (105.7 million miles)

Earth-Mars distance on landing day: 9:46 minutes
One-way speed-of-light time Mars-to-Earth on landing day: 500 million kilometers (311 million miles)
Total distance traveled Earth to Mars (approximate): 90 Mars days, or "sols" (equivalent to 92 Earth days)

Primary mission: 90 Mars days, or "sols" (equivalent to 92 Earth days)

Quick Facts

LOS ANGELES PUBLIC LIBRARY

3 7244 1819 0174 6



ST

An Astrobiology Strategy for the

EXPLORATION

of

MARS

THE NATIONAL ACADEMIES™

Advisers to the Nation on Science, Engineering, and Medicine

The nation turns to the National Academies—National Academy of Sciences, National Academy of Engineering, Institute of Medicine, and National Research Council—for independent, objective advice on issues that affect people's lives worldwide.

www.national-academies.org

ISBN-13 978-0-309-10851-5
ISBN-10 0-309-10851-9
9 0000
9 780309 108515

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

ST

PEARSON OPEN SOURCE SOFTWARE S

Linux® Hardening in Hostile Networks

Server Security from TLS to TOR

Kyle Rankin



Implement Industrial-Strength Security on Any Linux Server

In an age of mass surveillance, when advanced cyberwarfare weapons rapidly migrate into every hacker's toolkit, you can't rely on outdated security methods—especially if you're responsible for Internet-facing services. In *Linux® Hardening in Hostile Networks*, Kyle Rankin helps you to implement modern safeguards that provide maximum impact with minimum effort and to strip away old techniques that are no longer worth your time.

Rankin provides clear, concise guidance on modern workstation, server, and network hardening, and explains how to harden specific services, such as web servers, email, DNS, and databases. Along the way, he demystifies technologies once viewed as too complex or mysterious but now essential to mainstream Linux security. He also includes a full chapter on effective incident response that both DevOps and SecOps can use to write their own incident response plan.

Each chapter begins with techniques any sysadmin can use quickly to protect against entry-level hackers and presents intermediate and advanced techniques to safeguard against sophisticated and knowledgeable attackers, perhaps even state actors. Throughout, you learn what each technique does, how it works, what it does and doesn't protect against, and whether it would be useful in your environment.

- Apply core security techniques including 2FA and strong passwords
- Protect admin workstations via lock screens, disk encryption, BIOS passwords, and other methods
- Use the security-focused Tails distribution as a quick path to a hardened workstation
- Compartmentalize workstation tasks into VMs with varying levels of trust
- Harden servers with SSH, use apparmor and sudo to limit the damage attackers can do, and set up remote syslog servers to track their actions
- Establish secure VPNs with OpenVPN, and leverage SSH to tunnel traffic when VPNs can't be used
- Configure a software load balancer to terminate SSL/TLS connections and initiate new ones downstream
- Set up standalone Tor services and hidden Tor services and relays
- Secure Apache and Nginx web servers, and take full advantage of HTTPS
- Perform advanced web server hardening with HTTPS forward secrecy and ModSecurity web application firewalls
- Strengthen email security with SMTP relay authentication, SMTPS, SPF records, DKIM, and DMARC
- Harden DNS servers, deter their use in DDoS attacks, and fully implement DNSSEC
- Systematically protect databases via network access control, TLS traffic encryption, and encrypted data storage
- Respond to a compromised server, collect evidence, and prevent future attacks

Kyle Rankin, vice president of engineering operations for Final, Inc., is author of *DevOps Troubleshooting*, *The Official Ubuntu Server Book*, *Knoppix Hacks*, *Knoppix Pocket Reference*, *Linux Multimedia Hacks*, and *Ubuntu Hacks*, and is a contributor to several other books. An award-winning columnist for *Linux® Journal*, he has written for *PC Magazine*, *TechTarget*, and other media. His presentations on open source software include a keynote at SCALE 11x and numerous other talks at SCALE, O'Reilly Security Conference, OSCON, CactusCon, Linux World Expo, PenguinCon, and Linux users' groups.

Register Your Product ►►

at informit.com/register for convenient access to downloads, updates, and corrections as they become available.

ISBN-13: 978-0-13-417326-9

ISBN-10: 0-13-417326-0



9 780134 173269

informit.com/series/opensource

Cover design: Chuti Prasertsith
Cover photo: © VAlex/Shutterstock

Text printed on recycled paper



Pearson

Addison-Wesley

\$39.99 U.S. • \$49.99 CANADA

2nd Edition
Covers 802.11a, g, n & i

Creating & Administering Wireless Networks

ST

802.11® Wireless Networks

The Definitive Guide

O'REILLY®

Matthew S. Gast

Introduction to Wireless Networking

Over the past five years, the world has become increasingly mobile. As a result, traditional ways of networking the world have proven inadequate to meet the challenges posed by our new collective lifestyle. If users must be connected to a network by physical cables, their movement is dramatically reduced. Wireless connectivity, however, poses no such restriction and allows a great deal more freedom on the part of the network user. As a result, wireless technologies are encroaching on the traditional realm of “fixed” or “wired” networks. This change is obvious to anybody who drives on a regular basis. One of the “life and death” challenges to those of us who drive on a regular basis is the daily gauntlet of erratically driven cars containing mobile phone users in the driver’s seat.

Wireless connectivity for voice telephony has created a whole new industry. Adding mobile connectivity into the mix for telephony has had profound influences on the business of delivering voice calls because callers could be connected to people, not devices. We are on the cusp of an equally profound change in computer networking. Wireless telephony has been successful because it enables people to connect with each other regardless of location. New technologies targeted at computer networks promise to do the same for Internet connectivity. The most successful wireless data networking technology this far has been 802.11.

In the first edition of this book, I wrote about 802.11 being the tip of the trend in mobile data networking. At the time, 802.11 and third-generation mobile technologies were duking it out for mindshare, but 802.11 has unquestionably been more successful to date.

Why Wireless?

To dive into a specific technology at this point is getting a bit ahead of the story, though. Wireless networks share several important advantages, no matter how the protocols are designed, or even what type of data they carry.

The most obvious advantage of wireless networking is *mobility*. Wireless network users can connect to existing networks and are then allowed to roam freely. A mobile

telephone user can drive miles in the course of a single conversation because the phone connects the user through cell towers. Initially, mobile telephony was expensive. Costs restricted its use to highly mobile professionals such as sales managers and important executive decision makers who might need to be reached at a moment's notice regardless of their location. Mobile telephony has proven to be a useful service, however, and now it is relatively common in the United States and extremely common among Europeans.*

Likewise, wireless data networks free software developers from the tethers of an Ethernet cable at a desk. Developers can work in the library, in a conference room, in the parking lot, or even in the coffee house across the street. As long as the wireless users remain within the range of the base station, they can take advantage of the network. Commonly available equipment can easily cover a corporate campus; with some work, more exotic equipment, and favorable terrain, you can extend the range of an 802.11 network up to a few miles.

Wireless networks typically have a great deal of *flexibility*, which can translate into rapid deployment. Wireless networks use a number of base stations to connect users to an existing network. (In an 802.11 network, the base stations are called *access points*.) The infrastructure side of a wireless network, however, is qualitatively the same whether you are connecting one user or a million users. To offer service in a given area, you need base stations and antennas in place. Once that infrastructure is built, however, adding a user to a wireless network is mostly a matter of authorization. With the infrastructure built, it must be configured to recognize and offer services to the new users, but authorization does not require more infrastructure. Adding a user to a wireless network is a matter of configuring the infrastructure, but it does not involve running cables, punching down terminals, and patching in a new jack.[†]

Flexibility is an important attribute for service providers. One of the markets that many 802.11 equipment vendors have been chasing is the so-called “hot spot” connectivity market. Airports and train stations are likely to have itinerant business travelers interested in network access during connection delays. Coffeehouses and other public gathering spots are social venues in which network access is desirable. Many cafes already offer Internet access; offering Internet access over a wireless network is a natural extension of the existing Internet connectivity. While it is possible to serve a fluid group of users with Ethernet jacks, supplying access over a wired network is problematic for several reasons. Running cables is time-consuming and expensive

* While most of my colleagues, acquaintances, and family in the U.S. have mobile telephones, it is still possible to be a holdout. In Europe, it seems as if everybody has a mobile phone—one cab driver in Finland I spoke with while writing the first edition of this book took great pride in the fact that his family of four had six mobile telephones!

† This simple example ignores the challenges of scale. Naturally, if the new users will overload the existing infrastructure, the infrastructure itself will need to be beefed up. Infrastructure expansion can be expensive and time-consuming, especially if it involves legal and regulatory approval. However, my basic point holds: adding a user to a wireless network can often be reduced to a matter of configuration (moving or changing bits) while adding a user to a fixed network requires making physical connections (moving atoms), and moving bits is easier than moving atoms.

and may also require construction. Properly guessing the correct number of cable drops is more an art than a science. With a wireless network, though, there is no need to suffer through construction or make educated (or wild) guesses about demand. A simple wired infrastructure connects to the Internet, and then the wireless network can accommodate as many users as needed. Although wireless LANs have somewhat limited bandwidth, the limiting factor in networking a small hot spot is likely to be the cost of WAN bandwidth to the supporting infrastructure.

Flexibility may be particularly important in older buildings because it reduces the need for construction. Once a building is declared historical, remodeling can be particularly difficult. In addition to meeting owner requirements, historical preservation agencies must be satisfied that new construction is not desecrating the past. Wireless networks can be deployed extremely rapidly in such environments because there is only a small wired network to install.

Flexibility has also led to the development of grassroots community networks. With the rapid price erosion of 802.11 equipment, bands of volunteers are setting up shared wireless networks open to visitors. Community networks are also extending the range of Internet access past the limitations for DSL into communities where high-speed Internet access has been only a dream. Community networks have been particularly successful in out-of-the way places that are too rugged for traditional wireline approaches.

Like all networks, wireless networks transmit data over a network medium. The medium is a form of electromagnetic radiation.¹ To be well-suited for use on mobile networks, the medium must be able to cover a wide area so clients can move throughout a coverage area. Early wireless networks used infrared light. However, infrared light has limitations; it is easily blocked by walls, partitions, and other office construction. Radio waves can penetrate most office obstructions and offer a wider coverage range. It is no surprise that most, if not all, 802.11 products on the market use the radio wave physical layer.

Radio Spectrum: The Key Resource

Wireless devices are constrained to operate in a certain frequency band. Each band has an associated *bandwidth*, which is simply the amount of frequency space in the band. Bandwidth has acquired a connotation of being a measure of the data capacity of a link. A great deal of mathematics, information theory, and signal processing can be used to show that higher-bandwidth slices can be used to transmit more information. As an example, an analog mobile telephony channel requires a 20-kHz bandwidth. TV signals are vastly more complex and have a correspondingly larger bandwidth of 6 MHz.

¹ Laser light is also used by some wireless networking applications, but the extreme focus of a laser beam makes it suited only for applications in which the ends are stationary. "Fixed wireless" applications, in which lasers replace other access technology such as leased telephone circuits, are a common application.

infrared light has limitations; it is easily blocked by walls, partitions, and other office construction. Radio waves can penetrate most office obstructions and offer a wider coverage range. It is no surprise that most, if not all, 802.11 products on the market use the radio wave physical layer.

Radio Spectrum: The Key Resource

Wireless devices are constrained to operate in a certain frequency band. Each band has an associated *bandwidth*, which is simply the amount of frequency space in the band. Bandwidth has acquired a connotation of being a measure of the data capacity of a link. A great deal of mathematics, information theory, and signal processing can be used to show that higher-bandwidth slices can be used to transmit more information. As an example, an analog mobile telephony channel requires a 20-kHz bandwidth. TV signals are vastly more complex and have a correspondingly larger bandwidth of 6 MHz.

- Laser light is also used by some wireless networking applications, but the extreme focus of a laser beam makes it suited only for applications in which the ends are stationary. “Fixed wireless” applications, in which lasers replace other access technology such as leased telephone circuits, are a common application.

Early Adoption of 802.11

802.11's explosive advance has not been even. Some markets have evolved more quickly than others because the value of wireless networks is more pronounced in some markets. In general, the higher the value placed on mobility and flexibility, the greater the interest in wireless LANs.

Logistics organizations responsible for moving goods around (think UPS, FedEx, or airlines), were perhaps the earliest adopters of 802.11. Well before the advent of 802.11, package tracking was done with proprietary wireless LANs. Standardized products lowered the price and enabled competition between suppliers of network equipment, and it was an easy decision to replace proprietary products with standardized ones.

Health care has been an early adopter of wireless networks because of the great flexibility that is often required of health care equipment. Patients can be moved throughout a hospital, and the health care professionals that spend time with patients are among some of the most mobile workers in the economy. Technologically advanced health care organizations have adopted wireless LANs to make patient information available over wireless LANs to improve patient care by making information more accessible to doctors. Computerized records can be transferred between departments without the requirement to decipher the legendarily illegible doctor scrawls. In the cluttered environments of an emergency room, rapid access to imaging data can quite literally be a lifesaver. Several hospitals have deployed PCs to make radiology images available over wireless LANs on specially-equipped "crash carts" that offer instant access to X-rays, allowing doctors to make quick decisions without waiting for film to be developed.

Many educational institutions have enthusiastically adopted wireless LANs. 10 years ago, colleges competed for students based on how "wired" the campus was. More high speed data ports everywhere was assumed to be better. Nowadays, the leading stories in education are the colleges using wireless LANs to blanket coverage throughout the campus. Students are highly mobile network users, and can benefit greatly from network access between classes or in their "homes away from home" (the library, studio, or science lab, depending on major).

Radio spectrum allocation is rigorously controlled by regulatory authorities through licensing processes. Most countries have their own regulatory bodies, though regional regulators do exist. In the U.S., regulation is done by the Federal Communications Commission (FCC). Many FCC rules are adopted by other countries throughout the Americas. European allocation is performed by the European Radiocommunications Office (ERO). Other allocation work is done by the International Telecommunications Union (ITU). To prevent overlapping uses of the radio waves, frequency is allocated in bands, which are simply ranges of frequencies available to specified applications. Table 1-1 lists some common frequency bands used in the U.S.*

* The full spectrum allocation map is available from the National Telecommunications and Information Administration at <http://www.ntia.doc.gov/osmhome/allocchart.pdf>.

Table 1-1. Common U.S. frequency bands

Band	Frequency range
UHF ISM	902–928 MHz
S-Band	2–4 GHz
S-Band ISM	2.4–2.5 GHz
C-Band	4–8 GHz
C-Band satellite downlink	3.7–4.2 GHz
C-Band Radar (weather)	5.25–5.925 GHz
C-Band ISM	5.725–5.875 GHz
C-Band satellite uplink	5.925–6.425 GHz
X-Band	8–12 GHz
X-Band Radar (police/weather)	8.5–10.55 GHz
Ku-Band	12–18 GHz
Ku-Band Radar (police)	13.4–14 GHz 15.7–17.7 GHz

The ISM bands

In Table 1-1, there are three bands labeled ISM, which is an abbreviation for *industrial, scientific, and medical*. ISM bands are set aside for equipment that, broadly speaking, is related to industrial or scientific processes or is used by medical equipment. Perhaps the most familiar ISM-band device is the microwave oven, which operates in the 2.4-GHz ISM band because electromagnetic radiation at that frequency is particularly effective for heating water.

I pay special attention to the ISM bands in the table because those bands allow license-free operation, provided the devices comply with power constraints. 802.11 operates in the ISM bands, along with many other devices. Common cordless phones operate in the ISM bands as well. 802.11b and 802.11g devices operate within the 2.4 GHz ISM band, while 802.11a devices operate in the 5 GHz band.

The more common 802.11b/g devices operate in S-band ISM. The ISM bands are generally license-free, provided that devices are low-power. How much sense does it make to require a license for microwave ovens, after all? Likewise, you don't need a license to set up and operate a low-power wireless LAN.

What Makes Wireless Networks Different

Wireless networks are an excellent complement to fixed networks, but they are not a replacement technology. Just as mobile telephones complement fixed-line telephony, wireless LANs complement existing fixed networks by providing mobility to users. Servers and other data center equipment must access data, but the physical location of the server is irrelevant. As long as the servers do not move, they may as well be

connected to wires that do not move. At the other end of the spectrum, wireless networks must be designed to cover large areas to accommodate fast-moving clients. Typical 802.11 access points do not cover large areas, and would have a hard time coping with users on rapidly-moving vehicles.

Lack of Physical Boundary

Traditional network security places a great deal of emphasis on physical security of the network components. Data on the network travels over well-defined pathways, usually of copper or fiber, and the network infrastructure is protected by strong physical access control. Equipment is safely locked away in wiring closets, and set up so that it cannot be reconfigured by users. Basic security stems from the (admittedly marginal) security of the physical layer. Although it is possible to tap or redirect signals, physical access control makes it much harder for an intruder to gain surreptitious access to the network.

Wireless networks have a much more open network medium. By definition, the network medium in a wireless network is not a well-defined path consisting of a physical cable, but a radio link with a particular encoding and modulation. Signals can be sent or received by anybody in possession of the radio techniques, which are of course well known because they are open standards. Interception of data is child's play, given that the medium is open to anybody with the right network interface, and the network interface can be purchased for less than \$50 at your local consumer electronics store. Careful shopping online may get you cards for half of that.

Furthermore, radio waves tend to travel outside their intended location. There is no abrupt physical boundary of the network medium, and the range at which transmissions can be received can be extended with high-gain antennas on either side. When building a wireless network, you must carefully consider how to secure the connection to prevent unauthorized use, traffic injection, and traffic analysis. With the maturation of wireless protocols, the tools to authenticate wireless users and properly encrypt traffic are now well within reach.

Dynamic Physical Medium

Once a wired network is put in place, it tends to be boring, which is to say, predictable. Once the cables have been put in place, they tend to do the same thing day in and day out. Provided the network has been designed according to the engineering rules laid out in the specification, the network should function as expected. Capacity can be added to a wired network easily by upgrading the switches in the wiring closet.

In contrast, the physical medium on wireless LANs is much more dynamic. Radio waves bounce off objects, penetrate through walls, and can often behave somewhat unpredictably. Radio waves can suffer from a number of propagation problems that may interrupt the radio link, such as multipath interference and shadows. Without a

reliable network medium, wireless networks must carefully validate received frames to guard against frame loss. Positive acknowledgment, the tactic used by 802.11, does an excellent job at assuring delivery at some cost to throughput.

Radio links are subject to several additional constraints that fixed networks are not. Because radio spectrum is a relatively scarce resource, it is carefully regulated. Two ways exist to make radio networks go faster. Either more spectrum can be allocated, or the encoding on the link can be made more sensitive so that it packs more data in per unit of time. Additional spectrum allocations are relatively rare, especially for license-free networks. 802.11 networks have kept the bandwidth of a station's radio channel to approximately 30 MHz, while developing vastly improved encoding to improve the speed. Faster coding methods can increase the speed, but do have one potential drawback. Because the faster coding method depends on the receiver to pick out subtle signal differences, much greater signal-to-noise ratios are required. Higher data rates therefore require the station to be located closer to its access point. Table 1-2 shows the standardized physical layers in 802.11 and their respective speeds.

Table 1-2. Comparison of 802.11 physical layers (PHYs)

IEEE standard	Speed	Frequency band	Notes
802.11	1 Mbps	2.4 GHz	First PHY standard (1997). Featured both frequency-hopping and direct-sequence modulation techniques.
	2 Mbps		
802.11a	Up to 54 Mbps	5 GHz	Second PHY standard (1999), but products not released until late 2000.
802.11b	5.5 Mbps 11 Mbps	2.4 GHz	Third PHY standard, but second wave of products. The most common 802.11 equipment as the first edition of this book was written, and the majority of the legacy installed base at the time the second edition was written.
802.11g	Up to 54 Mbps	2.4 GHz	Fourth PHY standard (2003). Applies the coding techniques of 802.11a for higher speed in the 2.4 GHz band, while retaining backwards compatibility with existing 802.11b networks. The most common technology included with laptops in 2005.

Radio is inherently a broadcast medium. When one station transmits, all other stations must listen. Access points act much like old shared Ethernet hubs in that there is a fixed amount of transmission capacity per access point, and it must be shared by all the attached users. Adding capacity requires that the network administrator add access points while simultaneously reducing the coverage area of existing access points.

Security

Many wireless networks are based on radio waves, which makes the network medium inherently open to interception. Properly protecting radio transmissions on any network is always a concern for protocol designers. 802.11 did not build in much in the way of security protocols. Coping with the inherent unreliability of the wireless medium and mobility required several protocol features to confirm frame

delivery, save power, and offer mobility. Security was quite far down the list, and proved inadequate in the early specifications.

Wireless networks must be strongly authenticated to prevent use by unauthorized users, and authenticated connections must be strongly encrypted to prevent traffic interception and injection by unauthorized parties. Technologies that offer strong encryption and authentication have emerged since the first edition of this book, and are a major component of the revisions for the second edition.

A Network by Any Other Name...

Wireless networking is a hot industry segment. Several wireless technologies have been targeted primarily for data transmission. Bluetooth is a standard used to build small networks between peripherals: a form of “wireless wires,” if you will. Most people in the industry are familiar with the hype surrounding Bluetooth, though it seems to have died down as real devices have been brought to market. In the first edition, I wrote that I have not met many people who have used Bluetooth devices, but it is much more common these days. (I use a Bluetooth headset on a regular basis.)

Post-second-generation (2.5G) and third-generation (3G) mobile telephony networks are also a familiar wireless technology. They promise data rates of megabits per cell, as well as the “always on” connections that have proven to be quite valuable to DSL and cable modem customers. After many years of hype and press from 3G equipment vendors, the rollout of commercial 3G services is finally underway. 2.5G services like GPRS, EDGE, and 1xRTT are now widely available, and third-generation networks based on UMTS or EV-DO are quickly being built. (I recently subscribed to an unlimited GPRS service to get connected during my train trips between my office and my home.) Many articles quote peak speeds for these technologies in the hundreds of kilobits per second or even megabits, but this capacity must be shared between all users in a cell. Real-world downstream speeds are roughly comparable to dial-up modem connections and cannot touch an 802.11 hot spot.

This is a book about 802.11 networks. 802.11 goes by a variety of names, depending on who is talking about it. Some people call 802.11 *wireless Ethernet*, to emphasize its shared lineage with the traditional wired Ethernet (802.3). A second name which has grown dramatically in popularity since the first edition of this book is *Wi-Fi*, from the interoperability certification program run by the Wi-Fi Alliance, the major trade association of 802.11 equipment vendors. The Wi-Fi Alliance, formerly known as the Wireless Ethernet Compatibility Alliance (WECA), will test member products for compatibility with 802.11 standards.¹ Other organizations will perform compati-

¹ More details on the Wi-Fi Alliance and its certification program can be found at <http://www.wi-fi.org/>.

802.11 Wireless Networks: The Definitive Guide



Using a wireless network is a liberating experience. But underneath the experience lies a complex protocol, and even more complex issues arise when your data isn't limited to traveling on physical wires. How do you structure your network so mobile users can move around effectively? How do you extend wireless coverage so it's available everywhere you need it? What kinds of security issues do wireless networks raise? How do you tune your network for optimal performance? How do you provide enough capacity to support the users you expect initially, and how do you deal with the problems that arise as more users join the network?

802.11 Wireless Networks: The Definitive Guide, Second Edition discusses all these issues, and more. This book is for the serious system or network administrator who is responsible for deploying or maintaining a wireless network. It contains an extensive discussion of wireless security issues, including the problems with the WEP standard and a look at the alternatives: dynamic WEP, plus the 802.1X and the 802.11i security standards. Since network monitoring is essential to any serious network administrator, a chapter is devoted to network analysis and troubleshooting, using Ethereal and other tools.

802.11 Wireless Networks: The Definitive Guide brings you up-to-date on all the latest developments in wireless networking. In addition to 802.11b and 11a, this edition covers 802.11g and looks ahead to the 802.11n protocol, which is currently being standardized. This new edition greatly expands the discussion of network planning and architecture, paying special attention to mobility between access points, spectrum management, and power control. It's the only book available that discusses how to calculate the performance of your wireless network and how to tune your network for optimal performance.

Finally, *802.11 Wireless Networks: The Definitive Guide* shows you how to configure wireless cards and Linux, Windows, and Mac OS X systems, and how to work with access points. Few books in any field combine the theory you need to know with the practical experience and advice you need to get things working. *802.11 Wireless Networks: The Definitive Guide* is one of those books. If you are responsible for a wireless network, you need this book.

Matthew S. Gast is the leading author on the planning and deployment of wireless networks.

www.oreilly.com

ISBN 0-596-10052-3



US \$44.95

CAN \$62.95



9 780596 100520

Safari
BOOKS ONLINE
ENABLED

Includes
FREE 45-Day
Online Edition

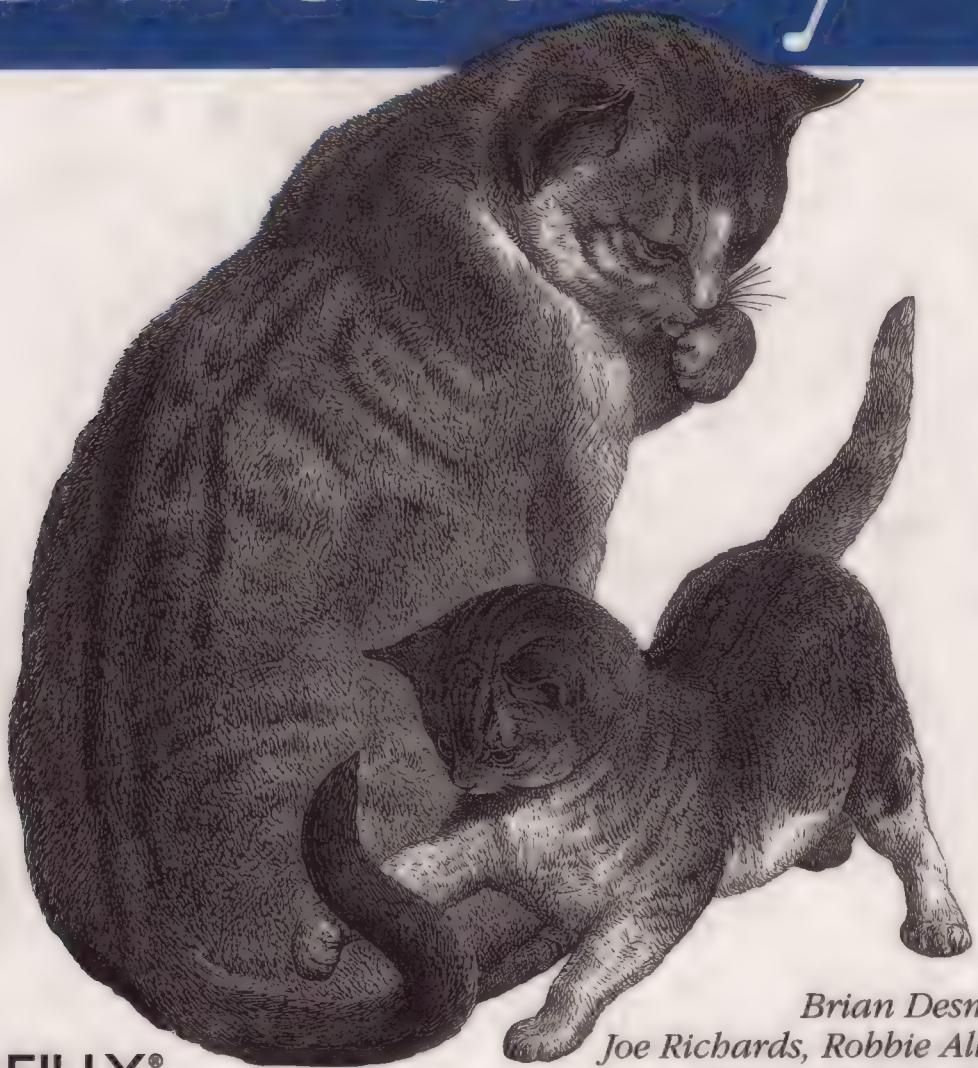
*Designing, Deploying, and
Running Active Directory*

ST

Cover
2003 - 2012

Windows Server
and PowerShell

Active Directory



Brian Desmond,
Joe Richards, Robbie Allen &
Alistair G. Lowe-Norris

O'REILLY®

Win'
DURALAM
PRO PART

LOS ANGELES PUBLIC LIBRARY



3 7244 2104 7199 2

Active Directory

Organize your network resources by learning how to design, manage, and maintain Active Directory. Updated to cover Windows Server 2012, the fifth edition of this bestselling book gives you a thorough grounding in Microsoft's network directory service by explaining concepts in an easy-to-understand, narrative style.

You'll negotiate a maze of technologies for deploying a scalable and reliable AD infrastructure, with new chapters on management tools, searching the AD database, authentication and security protocols, and Active Directory Federation Services (ADFS). This book provides real-world scenarios that let you apply what you've learned—ideal whether you're a network administrator for a small business or a multinational enterprise.

- Upgrade Active Directory to Windows Server 2012
- Learn the fundamentals, including how AD stores objects
- Use the AD Administrative Center and PowerShell
- Learn AD Federation Services
- Search and gather AD data, using the LDAP query syntax
- Understand how Group Policy functions
- Design a new Active Directory forest
- Examine the Kerberos security protocol
- Get a detailed look at the AD replication process
- Explore AD Lightweight Directory Services

Brian Desmond, Microsoft MVP since 2003, is a consultant focused on Active Directory and identity management for some of the world's largest companies.

US \$54.99

CAN \$57.99

ISBN: 978-1-449-32002-7



5 5 4 9 9

9 781449 320027



Twitter: @oreillymedia
facebook.com/oreilly

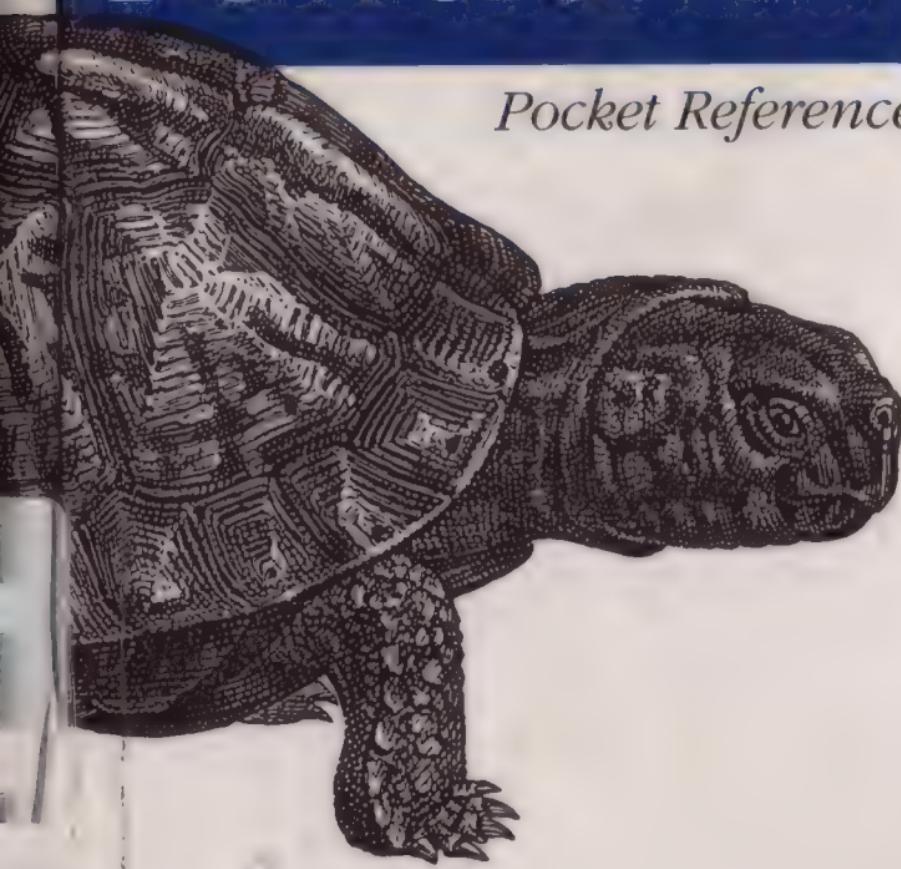
O'REILLY®
oreilly.com

*Portable Help for
PowerShell Scripters*

2nd Edition

Windows PowerShell

Pocket Reference



O'REILLY®

Lee Holmes

O'REILLY

ST

4th Edition
A GNU Manual

Effective awk Programming

UNIVERSAL TEXT PROCESSING AND PATTERN MATCHING



Arnold Robbins



3 7244 2210 7312 6

Effective awk Programming

When processing text files, the *awk* language is ideal for handling data extraction, reporting, and data-reformatting jobs. This practical guide serves as both a reference and tutorial for POSIX-standard *awk* and for the GNU implementation, called *gawk*. This book is useful for novices and *awk* experts alike.

In this thoroughly revised edition, author and *gawk* lead developer Arnold Robbins describes the *awk* language and *gawk* program in detail, shows you how to use *awk* and *gawk* for problem solving, and then dives into specific features of *gawk*. System administrators, programmers, webmasters, and other power users will find everything they need to know about *awk* and *gawk*. You will learn how to:

- Format text and use regular expressions in *awk* and *gawk*
- Process data using *awk*'s operators and built-in functions
- Manage data relationships using associative arrays
- Define your own functions
- "Think in *awk*" with two full chapters of sample functions and programs
- Take advantage of *gawk*'s many advanced features
- Debug *awk* programs with the *gawk* built-in debugger

This book is published under the terms of the GNU Free Documentation License. *You have the freedom to copy and modify this GNU manual.*

Royalties from the sales of this book go to the Free Software Foundation and to the author.

Arnold Robbins, a professional programmer and technical author, has worked with Unix systems since 1980 and has used *awk* since 1987. Arnold is the maintainer of *gawk* and its documentation. As a member of the POSIX 1003.2 balloting group, he helped shape the POSIX standard for *awk*.

TEXT PROCESSING

US \$44.99

CAN \$51.99

ISBN: 978-1-491-90461-9



5 4 4 9 9

9 781491 904619



Twitter: @oreillymedia
facebook.com/oreilly

Implement Industrial-Strength Security on Any Linux Server

In an age of mass surveillance, when advanced cyberwarfare weapons rapidly migrate into every hacker's toolkit, you can't rely on outdated security methods—especially if you're responsible for Internet-facing services. In *Linux® Hardening in Hostile Networks*, Kyle Rankin helps you to implement modern safeguards that provide maximum impact with minimum effort and to strip away old techniques that are no longer worth your time.

Rankin provides clear, concise guidance on modern workstation, server, and network hardening, and explains how to harden specific services, such as web servers, email, DNS, and databases. Along the way, he demystifies technologies once viewed as too complex or mysterious but now essential to mainstream Linux security. He also includes a full chapter on effective incident response that both DevOps and SecOps can use to write their own incident response plan.

Each chapter begins with techniques any sysadmin can use quickly to protect against entry-level hackers and presents intermediate and advanced techniques to safeguard against sophisticated and knowledgeable attackers, perhaps even state actors. Throughout, you learn what each technique does, how it works, what it does and doesn't protect against, and whether it would be useful in your environment.

Apply core security techniques including 2FA and strong passwords
Protect admin workstations via lock screens, disk encryption, BIOS
passwords, and other methods

Use the security-focused Tails distribution as a quick path to a hardened
workstation

Compartmentalize workstation tasks into VMs with varying levels of trust
Harden servers with SSH, use apparmor and sudo to limit the damage
attackers can do, and set up remote syslog servers to track their actions
Establish secure VPNs with OpenVPN, and leverage SSH to tunnel traffic
when VPNs can't be used

Configure a software load balancer to terminate SSL/TLS connections and
initiate new ones downstream

Set up standalone Tor services and hidden Tor services and relays

Secure Apache and Nginx web servers, and take full advantage of HTTPS

Perform advanced web server hardening with HTTPS forward secrecy and
ModSecurity web application firewalls

Strengthen email security with SMTP relay authentication, SMTPS, SPF
records, DKIM, and DMARC

Harden DNS servers, deter their use in DDoS attacks, and fully implement
DNSSEC

Systematically protect databases via network access control, TLS traffic
encryption, and encrypted data storage

Respond to a compromised server, collect evidence, and prevent future attacks

Linux Hardening Hostile Networks

Server Security from TLS to TOR

informat.com/series/opensource

Cover design: Chuti Prasertsith
Cover photo: © VAlex/Shutterstock

Text printed on recycled paper

ISBN-13: 978-0-13-417326-9
ISBN-10: 0-13-417326-0



5 3 9 9 9
9 7 8 0 1 3 4 1 7 3 2 6 9

\$39.99 U.S. • \$49.99 CANADA



Pearson
Addison-Wesley

Section 3

The third section of each chapter is where I have a bit of fun and go all out with advanced hardening steps aimed at advanced up to nation-state attackers. Some of these hardening steps are rather sophisticated and time-consuming, whereas others are really just the next step up from the intermediate approaches in Section 2. Although these steps are aimed at protecting against advanced threats, remember that today's advanced threats tend to find their way into tomorrow's script kiddie toolkits.

What This Book Covers

Now that we know how the chapters are structured, let's look at what each one covers.

Chapter 1: Overall Security Concepts

Before we get into specific hardening techniques, it's important to build a foundation with the security principles we will apply to all hardening techniques in the rest of the book. No security book can cover every possible type of threat or how to harden every type of application, but if you understand some of the basic concepts behind security you can apply them to whatever application you'd like to secure. Section 1 of Chapter 1 introduces some essential security concepts that you will apply throughout the book and finishes up with a section on choosing secure passwords and general password management. Section 2 elaborates on the security principles in the first section with a focus on more sophisticated attacks and provides a general introduction to two-factor authentication. Section 3 examines how general security principles apply in the face of an advanced attacker and discusses advanced password-cracking techniques.

Chapter 2: Workstation Security

A sysadmin workstation is a high-value target for an attacker or thief because administrators typically have privileged access to all servers in their environments. Chapter 2 covers a series of admin-focused workstation-hardening steps. Section 1 covers basic workstation-hardening techniques including the proper use of lock screens, suspend, and hibernation, and introduces the security-focused Linux distribution Tails as a quick path to a hardened workstation. The section finishes up by covering a few fundamental principles of how to browse the web securely including an introduction to HTTPS, concepts behind cookie security, and how to use a few security-enhancing browser plugins. Section 2 starts with a discussion of disk encryption, BIOS passwords, and other techniques to protect a workstation against theft, a nosy coworker, or a snooping customs official. The section also features more advanced uses of Tails as a high-security replacement for a traditional OS including the use of the persistent disk and the GPG clipboard applet. Section 3 covers advanced techniques such as using the Qubes OS to compartmentalize your different workstation tasks into their own VMs with varying levels of trust. With this in place if, for instance, your untrusted web browser VM gets compromised by visiting a bad website, that compromise won't put the rest of your VMs or your important files at risk.

Chapter 3: Server Security

If someone is going to compromise your server, the most likely attack will either be through a vulnerability in a web application or other service the server hosts, or through SSH. In other chapters, we cover hardening steps for common applications your server may host, so Chapter 3 focuses more on general techniques to secure just about any server you have, whether it's hosting a website, email, DNS, or something completely different. This chapter includes several techniques to harden SSH and covers how to limit the damage an attacker or even a malicious employee can do if he gains access to the server with tools like apparmor and sudo. We also cover disk encryption to protect data at rest and how to set up a remote syslog server to make it more difficult for an attacker to cover her tracks.

Chapter 4: Network

Along with workstation and server hardening, network hardening is a fundamental part of infrastructure security. Section 1 of Chapter 4 provides an overview of network security and then introduces the concept of the man-in-the-middle attack in the context of an attacker on an upstream network. Section 1 finishes up with an introduction to iptables firewall settings. Section 2 covers how to set up a secure private VPN using OpenVPN and how to leverage SSH to tunnel traffic securely when a VPN isn't an option. It then covers how to configure a software load balancer that can both terminate SSL/TLS connections and can initiate new ones downstream. Section 3 focuses on Tor servers, including how to set up a standalone Tor service strictly for internal use, as an external node that routes traffic within Tor, and as an external exit node that accepts traffic from the Internet. It also discusses the creation and use of hidden Tor services and how to set up and use hidden Tor relays for when you need to mask even that you are using Tor itself.

Chapter 5: Web Servers

Chapter 5 focuses on web server security and covers both the Apache and Nginx web servers in all examples. Section 1 covers the fundamentals of web server security including web server permissions and HTTP basic authentication. Section 2 discusses how to configure HTTPS, how to set it as the default by redirecting all HTTP traffic to HTTPS, how to secure HTTPS reverse proxies, and how to enable client certificate authentication. Section 3 discusses more advanced web server hardening including HTTPS forward secrecy and then web application firewalls with ModSecurity.

Chapter 6: Email

Email was one of the first services on the Internet, and it's still relied on by many people not just for communication but also security. Section 1 of Chapter 6 introduces overall email security fundamentals and server hardening, including how to avoid becoming an open relay. Section 2 covers how to require authentication for SMTP relays and how to enable SMTPS. Section 3 covers more advanced email security features that both aid in spam prevention and overall security such as SPF records, DKIM, and DMARC.

Chapter 7: DNS

Domain Name Service (DNS) is one of those fundamental network services to which many people never give a second thought (as long as it's working). In Chapter 7, we cover how to harden any DNS server before you put it on a network. Section 1 describes the fundamentals behind DNS security and how to set up a basic hardened DNS server. Section 2 goes into more advanced DNS features such as rate limiting to help prevent your server from being used in DDOS attacks, query logging to provide forensics data for your environment, and authenticated dynamic DNS. Section 3 provides an introduction to DNSSEC and the new DNSSEC records and discusses how to configure DNSSEC for your domain and how to set up and maintain DNSSEC keys.

Chapter 8: Database

If there is only one place in your infrastructure that holds important information, it's likely to be a database. In Chapter 8, we discuss a number of different approaches to database security for the two most popular open-source database servers: MySQL (MariaDB) and Postgres. Starting with Section 1, we cover some simple security practices you should follow as you set up your database. Section 2 then dives into some intermediate hardening steps including setting up network access control and encrypting traffic with TLS. Section 3 focuses on database encryption and highlights some of the options available for encrypted data storage in MySQL and Postgres.

Chapter 9: Incident Response

Even with the best intentions, practices, and efforts, sometimes an attacker still finds a way in. When that happens, you will want to collect evidence and try to find out how he got in and how to stop it from happening again. Chapter 9 covers how to best respond to a server you suspect is compromised, how to collect evidence, and how to use that evidence to figure out what the attacker did and how he got in. Section 1 lays down some fundamental guidelines for how to approach a compromised machine and safely shut it down so other parties can start an investigation. Section 2 gives an overview on how to perform your own investigation and discusses how to create archival images of a compromised server and how to use common forensics tools including Sleuth Kit and Autopsy to build a file system timeline to identify what the attacker did. Section 3 includes walking through an example investigation and guides to forensics data collection on cloud servers.

Appendix A: Tor

Chapter 4 discusses how to use Tor to protect your anonymity on a network, but it focuses more on how to use Tor and less about how Tor works. Here I dive a bit deeper into how Tor works and how it can protect your anonymity. I also discuss some of the security risks around Tor and how you can mitigate them.

Appendix B: SSL/TLS

Throughout the book, I explain how to protect various services with TLS. Instead of bogging you down with the details of how TLS works in almost every chapter, I've put those details here as a quick reference in case you are curious about how TLS works, how it protects you, its limitations, and some of its security risks and how to mitigate them.

Conventions

This book uses a monospace font for code. Code lines that exceed the width of the printed page are indicated by a continuation character (→) at the start of the portion of the line that has wrapped to indicate it is all one line.

Register your copy of *Linux® Hardening in Hostile Networks* on the InformIT site for convenient access to updates and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (9780134173269) and click Submit. Look on the Registered Products tab for an Access Bonus Content link next to this product, and follow that link to access any available bonus materials. If you would like to be notified of exclusive offers on new editions and updates, please check the box to receive email from us.

Incident Response

Even with the best intentions, practices, and efforts, sometimes an attacker still finds a way in. When that happens, you will want to collect evidence and try to find out how she got in and how to stop it from happening again. This chapter covers how to best respond to a server you suspect is compromised, how to collect evidence, and how to use that evidence to figure out what the attacker did and how she got in. “Section 1: Incident Response Fundamentals” lays down some fundamental guidelines for how to approach a compromised machine and safely shut it down so other parties can start an investigation. “Section 2: Secure Disk Imaging Techniques” gives an overview on how to perform your own investigation. It discusses how to create archival images of a compromised server and how to use common forensics tools including Sleuth Kit and Autopsy to build a file system timeline to identify what the attacker did. “Section 3: Walk Through a Sample Investigation” walks through an example investigation and guides to forensics data collection on cloud servers.

Section 1: Incident Response Fundamentals

Preparation before an attack occurs is just as important as the actions you take when it occurs. Even if you are naturally cool and calm during a crisis, there’s a good chance other members of your team won’t be, so a plan you have thought through when you are calm will be better than a plan you have thought up at the last minute with upper management breathing down your neck.

Who Performs Incident Response?

One important question to ask as you develop your incident response plan is, just who is responsible for incident response? In a small organization, the company may contract out incident response to a third party. In a large company, you may have an entire security operations center staffed with a team devoted solely to incident response. In either of these cases, your incident response plan may be simply to contact the primary party responsible for incident response. If, on the other hand, you are responsible for incident response, there are a number of additional policies you should work out ahead of time.

Do You Prosecute?

Before you develop any other specific responses, the first thing you should decide is under what circumstances you will wish to prosecute an attacker. If you are running a home

This command could take quite some time if the disk is large, but it's important to copy down the checksum and keep it somewhere safe so that you can prove later that no one has tampered with your working image.

Introduction to Sleuth Kit and Autopsy

The most simplistic way to start a forensics investigation would be to mount the images you have created as read-only on another machine and then simply look around the file system for evidence of an intrusion. As you get more sophisticated in your investigation, however, you'll find you want certain types of information repeatedly, such as the modified, access, and change times on files (MAC times); you'll want to keep track of checksums of any evidence you do find so you can prove later that the files weren't tampered with; you'll want to build an entire file system timeline that shows what files were changed or accessed in order of access; and you'll want to examine the file system for any deleted files the attacker may have created to cover their tracks. This is where tools like Sleuth Kit and Autopsy come in.

Sleuth Kit is a series of command-line tools aimed at forensics that make it easier to examine file system images. While Sleuth Kit provides a very powerful suite of tools, it also provides a steep learning curve to figure out all the correct command-line invocations you need to get the data you want. Autopsy works as a web-based front end to all the Sleuth Kit tools and makes it easy to examine a file system without learning each of the different command-line tools. Autopsy also makes it easy to organize multiple forensics analyses into different cases so you can reference them later. This makes all your Sleuth Kit commands more repeatable, protects you from some common mistakes during an investigation, and overall, keeps you organized.

Both Sleuth Kit and Autopsy should be packaged for most common Linux distributions, or you can download the software directly from <http://sleuthkit.org>. Note, however, that the Autopsy project has focused more on Windows with its latest revisions so we will discuss the last version of Autopsy that runs on Linux: version 2.

Once Autopsy and Sleuth Kit are installed, type `sudo autopsy` into a terminal to start the program. Autopsy needs root privileges so it can fully access block devices on your system as well as write to areas such as `/var/lib/autopsy` for evidence. Instructions on Autopsy's settings will appear in the terminal including the default location for evidence (`/var/lib/autopsy`) and the default port on which it listens (9999). Open a web browser and type in `http://localhost:9999/autopsy` to view the default Autopsy page and start your investigation (Figure 9-1).

From the main Autopsy page, click Open Case to open a case you have already created, or otherwise if you are starting from scratch (as we will be in this chapter) click New Case. The goal with a case is to organize all the disk images, investigators, and evidence for a specific attack in a single place so you can go back to it later. Each time you have a new compromised system or series of systems that are related to each other, you should start a new case.

On the New Case page, you can name and describe your case and you can also provide a list of investigators who will work on the case. Once your case is named

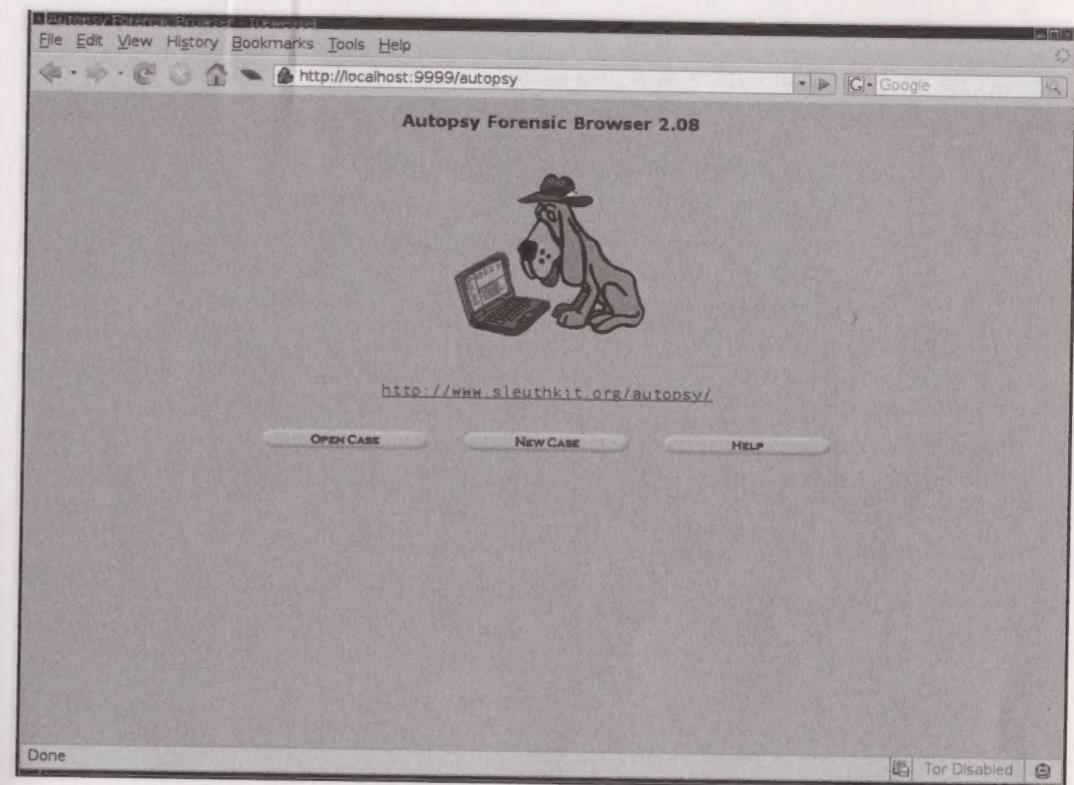


Figure 9-1 Default Autopsy page

and created, you will see the case gallery: a page that simply lists all the cases you have created. If this is your first case, just click OK to proceed to the Host Gallery. The Host Gallery lists all the servers you are investigating for this case. Often, an attacker will move from one compromised host to another, so include as many hosts as you need to investigate in this gallery. Like with the Case Gallery, click Add Host to fill out information about the host you are adding. You will see some interesting fields on the Add Host page relating to time. If the host was set to a time zone different from your local time zone, be sure to put its time zone in the Time Zone field. When you piece together a chain of events, especially across multiple hosts, having correctly synced time is valuable. The Timeskew Adjustment field lets you account for a server with out of sync time, and Autopsy will automatically adjust the times to reflect any skew you put in this field.

When you add the host and go back to the Host Gallery, select the host to analyze and click OK to go to the Host Manager page (Figure 9-2). If this is a new host, the first thing you should do is click Add Image File to add the image you have created previously. The image page only has three fields: Location, Type, and Import Method.

MARS

THE NASA MISSION REPORTS

